

LP-Based Robust Algorithms for Noisy Minor-Free and Bounded Treewidth Graphs

Nikhil Bansal*

Daniel Reichman†

Seeun William Umboh‡

Abstract

We give a general approach for solving optimization problems on noisy minor free graphs, where a δ -fraction of edges and vertices are adversarially corrupted. The noisy setting was first considered by Magen and Moharrami and they gave a $(1 + \delta)$ -estimation algorithm for the independent set problem. Later, Chan and Har-Peled designed a local search algorithm that finds a $(1 + O(\delta))$ -approximate independent set. However, nothing was known regarding other problems in the noisy setting. Our main contribution is a general LP-based framework that yields a $(1 + O(\delta \log m \log \log m))$ -approximation algorithm for noisy MAX- k -CSPs on m clauses.

1 Introduction

Several hard optimization problems often become substantially easier on special classes of graphs such as planar graphs and bounded treewidth graphs. For example, while the maximum independent set problem is notoriously hard on general graphs [27], it admits an efficient approximation scheme on planar graphs [32, 5] and can be solved exactly in polynomial time on bounded treewidth graphs [6]. Similarly, while MAX- k -SAT is APX-hard in general [28], planar instances admit an efficient approximation scheme [29] and bounded treewidth instances can be solved exactly in polynomial time [25, 35]. In general, there has been extensive work done on designing better algorithms for special graph classes and several general techniques have been developed for this purpose. For problems on bounded treewidth graphs, several techniques based on dynamic programming and deep results from algorithmic graph minor theory and logic have been developed [12, 6, 14, 11, 17]. For problems on planar graphs, many surprising approximation guarantees can be obtained based on decomposition approaches. One of the first decomposition approaches is based on the planar-separator theorem [32]. Later, Baker [5] developed a more versatile technique based on decomposition into $O(1)$ -outerplanar graphs (which have bounded treewidth). For example, Khanna and Motwani [29] used Baker’s technique to obtain efficient approximation schemes for a wide variety of Constraint Satisfaction Problems (CSPs) such as MAX-SAT.

Noisy graph models. In this paper, we consider a natural question that was first studied by Magen and Moharrami [33]: What happens to these special graph classes when they are perturbed adversarially? For the maximum independent set problem, [33] considers the setting where an

*Department of Mathematics and Computer Science, Eindhoven University of Technology, Netherlands. Email: n.bansal@tue.nl. Supported by NWO Vidi grant 639.022.211 and an ERC consolidator grant 617951.

†Electrical Engineering and Computer Science, University of California at Berkeley, Berkeley, CA 94709. Email: daniel.reichman@gmail.com.

‡Department of Mathematics and Computer Science, Eindhoven University of Technology, Netherlands. Email: seeun.umbogh@gmail.com. Supported by ERC consolidator grant 617951.

input graph G on n vertices is obtained from some (hidden) underlying planar graph G_0 by adding δn arbitrary edges (these are called *noisy* edges) for some small number $\delta > 0$ and ask: how well can one approximate the maximum independent set (MIS) problem on G ? More generally, one can consider the same question for other optimization problems that are easy on these special graph classes.

In this work, we consider MAX- k -SAT (for constant k) in the noisy setting. To introduce our noise model and to relate to the noisy graph model of [33], we remind the reader of the definition of a factor graph: given a k -SAT formula Φ , the *factor graph* of Φ is a bipartite graph $H = (A, B)$ where A contains a vertex for every variable appearing in Φ , B contains a vertex for every clause appearing in Φ , and a clause-vertex ϕ is connected to a variable-vertex x if and only if x belongs to the clause ϕ . In the noisy setting, the input formula Φ is given by an adversary who takes a planar k -SAT formula Φ_0 (i.e. Φ_0 's factor graph is planar) with n variables and m clauses, and adds δm clauses (each clause contains exactly k literals) to Φ_0 (resulting in δm vertices and $k\delta m = O(\delta m)$ edges being added to the factor graph of Φ_0). We choose to focus on this noise model as it does not change the arity of the original formula Φ_0 . Our results carry over without much difficulty to more general noise models where one adds both vertices and edges to the factor graph of Φ_0 , so long as the total number of edges and vertices added is δm .

Previous Work. Several previous works have considered approximation algorithms for MIS in the noisy setting. Magen and Moharrami presented an elegant argument showing that $\alpha(G)$ can be approximated to within a $(1+\epsilon)$ factor, for $\epsilon = \Omega(\delta)$, using $O(1/\epsilon)$ levels of the Sherali-Adams (SA) Hierarchy. Interestingly, this only yields an efficient *estimation* algorithm for $\alpha(G)$ and does not give any way to actually find the corresponding independent set.¹ In particular, the SA approach only uses the existence of G_0 to argue that $SA(G) \approx \alpha(G)$, without actually detecting the noisy edges. Later, Chan and Har-Peled [8] developed a PTAS for planar independent set based on local search which can be used to obtain a $(1+\epsilon)$ -approximation for MIS in noisy planar graphs. We are not aware of any previous work that studied approximation algorithms for noisy CSPs.

Limitations of Current Approaches. For MAX- k -SAT, classic approaches for planar instances based on computing decompositions of the factor graph break down completely in the presence of noise. In particular, the known algorithms for finding planar separators or decomposition into $O(1)$ -outerplanar graphs inherently use the planar structure in various ways, and are easily tricked even with very few noisy edges. In fact, a key motivation of [33] for studying the noisy setting was to design more “robust” algorithms that are not specifically tailor-made for particular graph classes. Note that in our noise model the adversary is even allowed to add a bounded degree expander on some subset of $O(\delta n)$ vertices, completely destroying the planar structure and increasing the treewidth to $\Omega(n)$. Another natural approach might be to recover some planar graph \tilde{G} from G without removing too many edges, and then apply the algorithms for planar graphs to \tilde{G} . However the best known guarantees for this Minimum Planarization problem are too weak for our purpose. In particular, even for bounded degree graphs these algorithms [10, 9] only achieve a $poly(n, \text{OPT})$ approximation, where OPT is the number of noisy edges, and thus only work in our setting when the noise parameter $\delta = n^{-\Omega(1)}$.

Finally, one may attempt to apply the local search algorithm of [8]. However, the analysis of the local search algorithm of [8] crucially relies on the existence of a decomposition based on the planar separator theorem known as r -divisions [24] and does not seem to apply to noisy instances of MAX- k -SAT.

¹Self-reducibility techniques such as those used for finding independent set in perfect graphs do not seem to work here as the estimation algorithm only gives an approximate answer.

1.1 Our Contributions

In this work, we give a general approach for solving problems on noisy planar graphs. Using this approach, we give the first algorithm that is able to handle noisy versions of planar MAX- k -SAT.

Theorem 1.1. *Let Φ_0 be a planar k -SAT formula over n variables and m clauses, where k is a constant independent of n and m . Let Φ be a k -SAT formula obtained by adding δm clauses to Φ_0 , for some $\delta > 0$. Then there is an algorithm such that given any $\epsilon > 0$, finds a $(1 + O(\epsilon + \delta \log m \log \log m))$ -approximate assignment in time $m^{O((\log \log m)^2/\epsilon)}$.*

Remark: Since MAX- k -SAT can be approximated within a constant factor, Theorem 1.1 is of interest when $\delta = O(1/(\log m \log \log m))$. Theorem 1.1 can also be extended to arbitrary (binary) k -CSPs.

For the maximum independent set problem, our approach gives an LP-based approximation algorithm with much better running time and whose approximation ratio has a better dependence on δ .

Theorem 1.2. *Let G be an n -vertex graph obtained by adding δn arbitrary edges to some planar graph G_0 , for some $\delta > 0$. Then there is an algorithm such that given any $\epsilon > 0$, finds an independent set of size within a $(1 + O(\epsilon + \delta))$ factor of $\alpha(G)$, and runs in time $n^{O(1/\epsilon^4)}$.*

General framework The starting point for our results is the following simple observation: Most algorithms for planar graphs use the planar structure only to find a certain structured decomposition of the graph, and once this is found, apply some simple or brute-force algorithm.

For example, the planar separator theorems are used to argue that given any $\epsilon > 0$, a planar graph can be decomposed into disjoint components of size $O(1/\epsilon^2)$ by removing some subset X of at most ϵn vertices. Similarly in Baker's decomposition, an ϵ fraction of edges F (or vertices) can be removed from a planar graph to decompose it into $O(1/\epsilon)$ -treewidth graphs.

Now consider a noisy planar graph G . We claim that such nice decompositions also exist for G (although it is unclear how to find them). For example, consider the decomposition of the underlying planar graph G_0 into bounded size pieces and for every noisy edge in $G \setminus G_0$ put one of its endpoint in X . This subset X has size $|X| \leq (\epsilon + \delta)n$ and removing X splits G into components of size $O(1/\epsilon^2)$. Similarly, let F be the edges removed in Baker's decomposition of G_0 , plus the noisy edges in $G \setminus G_0$. Clearly, $|F| \leq (\epsilon + \delta)n$ and removing F decomposes G into $O(1/\epsilon)$ -treewidth graphs.

So this leads to the natural question of whether we can directly find such good decompositions, without relying on the topological or other specific structure of the graph. Our main contribution is to show that this can indeed be done using general LP-based techniques. Specifically, we consider the following problems.

1. *Bounded Size Interdiction (weaker form):* Suppose G can be decomposed into components of size at most $1/\epsilon^2$ by removing some (small) subset X of vertices. Find such a subset X .
2. *Bounded Treewidth Interdiction:* Suppose G can be decomposed into graphs of treewidth at most w , by removing some subset F of edges. Find such a subset F .

Our Results. We show the following results for these general interdiction problems. Our main technical result is the following theorem.

Theorem 1.3. *Given a graph G and an integer $w > 0$, let F be some subset of edges such that removing them reduces the treewidth of G to w . Then there is an algorithm that runs in time $n^{O(1)}$ and finds a subset of edges F' such that $|F'| = O(\log n \log \log n)|F|$ and removing F' from G reduces the treewidth to $O(w \log w)$.*

Note that this gives a bicriteria $(\log n \log \log n, \log w)$ -approximation to the Bounded Treewidth Interdiction problem. We also remark that the approximation factors and the running time of the above algorithm do not depend on w .

Let us now see how Theorem 1.3 can be used to obtain Theorem 1.1. The basic idea is to apply Theorem 1.3 to the factor graph H of the noisy formula Φ to obtain a formula Φ' whose factor graph has a smaller treewidth, and then compute an exact solution for Φ' . Since k is constant, the noisy formula Φ has at most $O(m)$ variables and thus H has $\Theta(m)$ vertices. Thus, as discussed above, we know that there exists a set F of $(k\delta + O(\epsilon/(\log m \log \log m)))m$ edges that we can delete from H to reduce its treewidth to $O(\log m \log \log m/\epsilon)$. So we can apply Theorem 1.3 to H with $w = O(\log m \log \log m/\epsilon)$ to find an edge set F' of size at most $O((\epsilon + k\delta(\log m \log \log m))m)$ such that the residual factor graph $H - F'$ has treewidth $O(\log m(\log \log m)^2/\epsilon)$. Now observe that deleting the clauses incident to F' gives us a formula Φ' whose factor graph is a subgraph of $H - F'$, and thus has treewidth $O(\log m(\log \log m)^2/\epsilon)$. Thereafter, one can use an exact algorithm for MAX- k -SAT [29] that runs in time $2^{O(\ell)} \cdot \text{poly}(m)$ when the factor graph has treewidth at most ℓ . Next, we analyze the quality of this solution. Denote by OPT and OPT' the maximum number of satisfiable clauses in Φ and Φ' , respectively. At least a constant fraction of the clauses in Φ are satisfiable, so $\text{OPT} \geq \Omega(m)$. Since the formula Φ' is obtained by deleting $O(\epsilon + \delta(\log m \log \log m))m$ clauses from Φ (recall that k is a constant), we have $\text{OPT}' \geq \text{OPT} - O(\epsilon + \delta(\log m \log \log m))m$. Thus, $\text{OPT}' \geq (1 - O(\epsilon + \delta \log m \log \log m)) \text{OPT}$ and so our solution is a $(1 + O(\epsilon + \delta \log m \log \log m))$ -approximation.

We remark that Bounded Treewidth Interdiction is a fundamental problem of interest beyond its application to noisy planar graphs. For instance, treewidth interdiction algorithms have found applications in designing PTASes in several settings [22, 15, 23].

Next, we consider the Bounded Size Interdiction problem in Section 5.

Theorem 1.4. *(Weaker version) For the weaker form of the Bounded Size Interdiction problem stated above, given any $\beta \leq 1$, we can find in time $n^{O(1/\epsilon^2)}$ a subset of vertices X' with $|X'| \leq O(|X|/\beta + \beta|E|)$ such that $G[V \setminus X']$ has no component larger than $1/\epsilon^2$. Here, $|E|$ is the number of edges in G .*

For a noisy planar graph G , there exists $X \subseteq V$ of size $|X| \leq (\epsilon + \delta)n$ (as discussed above) and $|E| \leq (3 + \delta)n = O(1)n$. The latter follows as a planar graph has at most $3n - 6$ edges. So setting $\beta = (\epsilon + \delta)^{1/2}$ in Theorem 1.4 already gives X' of size $O((\epsilon + \delta)^{1/2}n)$ in time $n^{O(1/\epsilon^2)}$. This can be used to design a $(1 + O((\epsilon + \delta)^{1/2}))$ -approximation algorithm for independent set in noisy planar graphs. To get the better approximation factor of $1 + O(\epsilon + \delta)$ of Theorem 1.2 above, we will use a more refined result (Theorem 5.2) that decouples the dependence of $|X'|$ on ϵ and δ .

To prove Theorem 5.2, we write a configuration LP based formulation and round it suitably. (See Section 5.) The proof of Theorem 1.3 is much more challenging and requires several new ideas, and we give a broad overview of the algorithm and the proof below.

Overview of Techniques for Theorem 1.3: First, observe that if G has treewidth at most w , then $F = \emptyset$, and the algorithm must return $F' = \emptyset$. Thus, the problem is at least as hard as determining the treewidth of G . This is well known to be NP-Hard, and in fact unlikely to admit a polynomial time $O(1)$ approximation under reasonable complexity assumptions [40]. This

implies that the bicriteria guarantee is necessary, and that it is unlikely that the approximation with respect to w can be made $O(1)$.

At a high level, our algorithm will try to construct a good tree decomposition of width w , while removing some problematic edges along the way. To this end, let us first see how the known algorithms for finding tree decompositions work. Treewidth is characterized up to $O(1)$ factor by the well-linkedness property of a graph, which allows one to construct a tree decomposition by computing small balanced vertex separators recursively. Either the algorithm succeeds at each step and eventually finds a tree decomposition, or returns a well-linked set as a certificate that G has large treewidth. Finding balanced vertex separators is hard but one can use LP or SDP (resp.) formulations [7, 2, 20] based on spreading constraints [19, 4], and lose an $O(\log w)$ or $O(\sqrt{\log w})$ (resp.) factor in the quality of the treewidth.

In the noisy setting, our algorithmic task can thus be viewed as detecting which edges to remove so that the above recursive procedure works. To do this, we formulate an LP with variables for which edges to remove (let us call these the x_{uv} variables) so that in the residual graph every subset S of vertices has a small fractional balanced vertex separator of size at most w . However, as there are exponentially many such sets S , this gives a huge overall LP with (both) exponentially many variables and constraints, and it is unclear how to solve it. In particular, we have exponentially many different vertex separator LPs coupled together with the common x_{uv} variables.

We describe the algorithm in two parts. First, we assume that we are given the x_{uv} values from some feasible optimum LP solution. Using these x -values, for any given set S , we can now formulate a balanced edge-and-vertex separator LP, where the x -values give the fractional amount by which edges are removed, and in addition at most w vertices are removed. Using standard region-growing techniques jointly on these edge and vertex values, we decide which edges to delete (this adds to F'), and which vertices lie in the separator for S (these enter the bags in the tree decomposition). Doing this directly gives an $O(\log^2 n)$ approximation (provided we ensure that the separator tree is balanced and has depth $O(\log n)$), due to the loss of an $O(\log n)$ factor on each level of recursion. To reduce this to $O(\log n \log \log n)$, we use “Seymour’s trick” of more careful region-growing [38], together with some additional technical steps needed to make it work together with the tree decomposition procedure.

Second, we describe how to “solve” the LP. Perhaps surprisingly, this turns out to be quite challenging and requires some new ideas, which may be useful in other contexts. We only sketch these here, and details can be found in Section 4. First, we bypass the need to completely solve this LP, by using the Round-or-Separate framework (as in [3, 31]). In particular, the algorithm starts with some possibly infeasible solution x , and tries to construct the tree decomposition. If it succeeds, we are done. Otherwise, it gets stuck at finding a small balanced vertex separator for some set S . At that point, we try to add a violated inequality. However, a crucial point is that we need to find a violated inequality only involving the x -variables. So, a key step is to reformulate the LP to only have the x -variables. This crucially uses LP duality and the structure of the LP that the variables for different sets S are only loosely coupled via the x -variables. After this reformulation, it is still unclear how to find a violated inequality due to the exponential size of the LPs involved. We get around this issue by using some further properties of the Ellipsoid Method and the LP duality.

Other Related Work. The noise model considered here gives an interesting interpolation between easy and general worst-case instances. This is similar in spirit to approaches such as smoothed analysis [39], planted models and semi-random models [21, 34], although unlike these models our noise model is completely adversarial. For the general version of the bounded size interdiction

problem—given a graph G and a size parameter s , find the smallest vertex set X to delete so that $G - X$ has components of size at most s —Lee [30] independently gave a $O(\log s)$ -approximation algorithm that runs in time $2^{O(s)} \text{poly}(n)$. However, using this for maximum independent set on noisy planar graphs only yields a $(1 + O((\delta + \epsilon) \log(1/\epsilon)))$ -approximation. Fomin et al. [22] have considered the vertex deletion variant of our treewidth interdiction problem. They obtain a constant approximation factor for the problem, however their approximation factor depends at least exponentially on w which makes it inapplicable in our setting². Also their algorithm is polynomial time only for $w = O(1)$. A related model was considered by [26] in the context of property-testing and sublinear time algorithms in the bounded degree model.

2 Notation and Preliminaries

We always use G_0 for the underlying graph, and G for the noisy graph. The number of vertices of G is always n . For a subset $S \subseteq V$ and $F \subseteq E$, we use the notation $G[S] - F$ to denote the subgraph induced on the vertices S , excluding the edges in F . We use the notation $E(S)$ to denote the subset of E with both endpoints in S , and given another subset $S' \subseteq V$, we use $E(S, S')$ to denote the subset of E with one endpoint in S and another in S' .

Planar and Minor-Free graphs. The classic planar separator theorem [32] states that any planar graph has a $2/3$ -balanced vertex separator of size $O(\sqrt{n})$ and that it can be found efficiently. Applying this recursively gives the following.

Lemma 2.1. *For any planar graph G and any $\alpha > 0$, there is subset of vertices $X \subset V$ with $|X| = O(\alpha n)$, such that every component C_i of $G_0[V - X]$ has at most $1/\alpha^2$ vertices.*

A more generally applicable technique (see e.g. [29, 18, 13]) is Baker’s decomposition [5], which states that for any integer k , a planar graph can be decomposed into pieces of treewidth $O(k)$ (specifically, k -outerplanar graphs) by removing $O(1/k)$ fraction of edges or vertices.

A minor of G is a graph G' obtained by deleting and contracting edges. A graph G is H -free if G does not contain a subgraph H as a minor. Planar graphs are exactly the graphs excluding $K_{3,3}$ and K_5 as minors. In fact, Robertson and Seymour proved that every graph family closed under taking minors is characterized by a set of excluded minors. Both the planar separator theorem and Baker’s decomposition approach extend more generally to H -free graphs [1, 16, 15].

Treewidth. We review some relevant definitions related to treewidth.

Definition 2.2 (α -separator of S in G). *Given a graph $G = (V, E)$ and a set $S \subset V$, a vertex set $X \subset V$ is an α -separator of vertex set S in G if every component C of $G[V - X]$ has $|C \cap S| \leq \alpha|S|$.*

Definition 2.3 (Well-linked sets). *A vertex set S is w -linked in G if it does not have a $\frac{1}{2}$ -separator X with $|X| < w$. The linkedness of G is defined to be the maximum integer w such that there exists a w -linked set in G , and is denoted as $\text{link}(G)$.*

Definition 2.4 (Tree decomposition). *A tree decomposition of G is a tree T whose nodes t correspond to vertex subsets V_t of G (called bags) that satisfies the following properties: (i) for every edge $(u, v) \in E$, there exists a bag V_t containing both u and v ; (ii) for every vertex v , the bags that contain v form a non-empty subtree of T . The width of the decomposition is $\text{width}(T) = \max_{s \in T} |V_s| - 1$.*

²We need a sublinear dependence on w . To see why, consider for example the noisy MIS problem with G_0 as a grid. If we wish to reduce the treewidth to w , we would need to remove an $\Omega(1/w)$ fraction of the vertices, so if the interdiction algorithm is not an $o(w)$ approximation, it might end up deleting all the vertices.

$$\begin{array}{ll}
\min & \sum_{(u,v) \in E} x_{uv} \\
\text{s.t.} & \sum_{v \in V} y_v^S \leq w \quad \forall S \subseteq V \\
& d_{uv}^S \leq \sum_{e \in E(P)} x_e + \sum_{t \in V(P)} y_t^S \quad \forall S \subseteq V, u, v \in V, P \in \mathcal{P}(u, v) \\
& \sum_{v \in U} d_{uv}^S \geq |U| - \frac{|S|}{2} \quad \forall U \subseteq S \subseteq V, u \in U
\end{array} \tag{1}$$

Figure 1: LP relaxation for the treewidth interdiction problem.

Definition 2.5 (Treewidth). *The treewidth of G is the minimum integer w such that it has a tree decomposition of width w , and is denoted as $\text{tw}(G)$.*

The following well-known (see e.g. [36]) approximate characterization of treewidth in terms of linkedness will be useful for us.

Lemma 2.6 ([36]). *For any graph G , $\text{link}(G) < \text{tw}(G) < 4 \text{link}(G)$.*

3 Bounded Treewidth Interdiction

Recall that in the Bounded Treewidth Interdiction problem, we are given a graph $G = (V, E)$, a target treewidth w , and we want to find the minimum set F of edges to delete from G such that $\text{tw}(G - F) < w$. In this section, we describe the exponential-size LP and sketch the rounding algorithm used to prove Theorem 1.3. In the following, when X is a vertex set and F is an edge set, we use the shorthand $G - X$ to mean $G[V - X]$, and $G - X - F$ to mean $G[V - X] - F$.

3.1 An Exponential-Sized LP

Lemma 2.6 gives us a convenient characterization of feasible solutions F which we can use to write an LP. In particular, it says that if $\text{tw}(G - F) < w$, then every vertex set $S \subseteq V$ has a $\frac{1}{2}$ -separator X^S in $G - F$ of size less than w . Consider LP (1) in Figure 1. It has a variable x_{uv} indicating if edge $(u, v) \in E$ belongs to F . For every subset $S \subseteq V$ and vertex $v \in V$, variable y_v^S indicates if v belongs to the minimum-size $\frac{1}{2}$ -separator X^S of S in $G - F$. For $u, v \in V$, let $\mathcal{P}(u, v)$ denote the set of paths between u and v . For a path P , define $E(P)$ to be the set of edges in P and $V(P)$ to be the set of vertices on P , including the endpoints.

We interpret the solution as follows: the LP assigns a length x_{uv} to each edge $(u, v) \in E$ and a weight y_v^S for each vertex set S and vertex v . Consider a fixed set S . Without loss of generality, we can assume that the variable d_{uv}^S denotes the distance between u and v induced by the edge lengths x_e and vertex weights y_t^S . In particular, if we define the length of a path $P \in \mathcal{P}(u, v)$ to be the sum of edge lengths and vertex weights on the path, including the weights on u and v , then d_{uv}^S is the length of the shortest path between u and v . The variables d_{uv}^S and the last set of constraints are often called *spreading metrics* and *spreading constraints*, respectively. It is also easy to see that without loss of generality any feasible solution satisfies $d_{uv}^S \leq 1$. Note that there is a

potentially different metric d^S for each set S , and that the LP has exponentially many constraints and exponentially many variables.

Lemma 3.1. *LP (1) is a relaxation of the treewidth interdiction problem.*

Proof. We show that for every edge set F such that $\text{tw}(G - F) < w$, there exists a feasible solution $(\mathbf{x}, \mathbf{y}, \mathbf{d})$ to LP (1) with $\sum_{(u,v) \in E} x_{uv} \leq |F|$. Let \mathbf{x} be the indicator vector for F .

As $\text{tw}(G - F) < w$, by Lemma 2.6, for each vertex set S , there exists a set X^S of at most w vertices such that no component of $G - X^S - F$ contains more than half of S . Define \mathbf{y}^S to be the indicator vector for X^S and $d_{uv}^S = 1$ if either u or v lies in X^S , or if u and v lie in separate components of $G - X^S - F$.

The solution $(\mathbf{x}, \mathbf{y}, \mathbf{d})$ has $\sum_{(u,v) \in E} x_{uv} = |F|$ and satisfies the first two sets of constraints. It remains to show that it satisfies the spreading constraints. Fix some S and consider $U \subseteq S$ with $|U| > |S|/2$. Let u be a vertex of U . There are two cases to consider: (i) If $u \in X^S$, we have $d_{uv}^S = 1$ for all $v \in U \setminus \{u\}$ and so $\sum_{v \in U} d_{uv}^S \geq |U| - 1 \geq |U| - |S|/2$; (ii) otherwise if $u \notin X^S$, let C be the component of $G - X^S - F$ that contains u . We have $|C \cap S| \leq |S|/2$ since X^S is a $\frac{1}{2}$ -separator of S in $G - F$. We also have $d_{uv}^S = 1$ for every $v \in U - (C \cap S)$ since v is either in a different component of $G - X^S - F$ or in X^S . Thus,

$$\sum_{v \in U} d_{uv}^S \geq \sum_{v \in U - (C \cap S)} d_{uv}^S \geq |U| - |C \cap S| \geq |U| - |S|/2.$$

Therefore, the solution $(\mathbf{x}, \mathbf{y}, \mathbf{d})$ is a feasible solution with $\sum_{(u,v) \in E} x_{uv} = |F|$. \square

In the rest of this section, we describe our rounding algorithm for LP (1) and prove the following lemma.

Lemma 3.2. *Given oracle access to a feasible solution $(\mathbf{x}, \mathbf{y}, \mathbf{d})$ of LP (1), we can find in time $\text{poly}(n)$ an edge set F such that $\text{tw}(G - F) \leq O(w \log w)$ and $|F| \leq O(\log n \log \log n) \sum_{(u,v) \in E} x_{uv}$.*

3.2 Sketch of the Rounding Algorithm

Let $(\mathbf{x}, \mathbf{y}, \mathbf{d})$ be a feasible solution to LP (1). As mentioned in the Introduction, our rounding algorithm is based on a recursive algorithm for constructing tree decompositions with Seymour's recursive graph decomposition trick layered on top. We now describe at a high level how these ideas are combined together and highlight the key issues that arise.

Classic Tree Decomposition. We begin by outlining the relevant parts of the classic tree decomposition algorithm [37] (which we call DECOMPOSITION). It is a recursive algorithm: Given a subgraph H of G and an integer w , it either constructs a tree decomposition of width $O(w)$ or it finds a w -linked set S which certifies³ that $\text{tw}(G) \geq w$. There are two key steps (illustrated in Figure 2) that are important to us:

- **Separate:** find a minimum-size $\frac{2}{3}$ -separator X of a vertex set S in H (the particular choice of S depends on previous recursive steps).
- **Recurse:** for each component C_i of $H - X$, recurse on the subgraph H_i of H which consists of the edges of H induced by C_i and those between C_i and X .

³It certifies $\text{tw}(H) \geq w$, and the treewidth of a graph is at least the treewidth of any subgraph.




Figure 2: DECOMPOSITION finds a separator X , and for each component C_i of $H - X$, it recurses on the subgraph H_i consisting of edges of H with at least one endpoint in C_i .




Figure 3: INTERDICT uses the spreading metric given by LP solution (left), partitions the metric into regions with cut vertices shown as hollow vertices and cut edges shown as dashed edges (center), and for each region, recurses on the subgraph contained in it (right). Note that the subgraph includes the cut vertices on the boundary of the region as well.

We say that DECOMPOSITION “succeeds” if it does not encounter a w -linked vertex set S . In particular, when DECOMPOSITION succeeds, the separators found during its execution can be used to construct a tree decomposition of width $O(w)$; when it fails, it has found a w -linked set S during its recursion.

Our Algorithm. Our algorithm (which we call INTERDICT) largely follows along the lines of DECOMPOSITION. The main difference is that we also want to delete edges to ensure that size of the separators found in the recursion are small enough so that we succeed in constructing a tree decomposition of width $O(w \log w)$. We still make the same choices about which set S to separate and how to recurse; this allows us to reuse the analysis of DECOMPOSITION to prove that we have deleted enough edges to reduce the treewidth down to $O(w \log w)$. In particular, instead of the Separate step, we want to perform a “Delete and Separate” step instead.

- Delete and Separate: delete a subset D of edges and find a vertex set X of size $O(w \log w)$ such that X is a $\frac{2}{3}$ -separator of S in $H - D$.

Here is where LP (1) is useful. It is similar to the spreading metric relaxation for finding minimum balanced vertex separators, except that it gives edge-and-vertex separators. As mentioned in the Introduction, one can apply standard region growing techniques in an almost black box fashion along with other tricks to obtain a $(\log^2 n, \log w)$ -approximation to treewidth interdiction.

Obtaining a $(\log n \log \log n, \log w)$ -approximation needs some care. At a high level, we want to apply region growing recursively using Seymour’s recursion. The basic idea is to ensure that not only is $|D|$ bounded by the cost of the LP solution projected on H , but it is also bounded in some nice way by the cost of the LP solution on the subgraphs H_i we recurse on. In particular, we want to use region growing to partition the spreading metric d^S into “regions” B_1, \dots, B_k (illustrated in Figure 3) with the following properties:

- (bounded charge) the set of edges $\delta_x(B_i)$ (vertices $\delta_y(B_i)$ resp.) cut by B_i can be charged to the x_{uv} variables (y_v^S variables, resp.) “contained” in the region,
- (containment) each subgraph H_i we recurse on is contained inside a region.

We can then choose $D = \delta_x(B_1) \cup \dots \cup \delta_x(B_k)$ and $X = \delta_y(B_1) \cup \dots \cup \delta_y(B_k)$. Due to the structure of the subproblems in the Recurse step, ensuring the second property requires some care with how we implement region growing. The problem is that the separators involve both edges and vertices, and they play different roles, i.e. edges are deleted globally, while vertices are deleted locally for each S . Normally, the region growing technique proceeds by finding a region B_i in the current graph that satisfies the bounded charge property, removes $B_i \cup \delta_y(B_i)$ —the vertices contained in or cut by B_i from the graph—and repeats on the residual graph. However, this would also remove any edge (u, v) with $u \in \delta_y(B_i)$, even though v is still remaining in the residual graph. In this case, no future region can contain or cut the edge (u, v) . This is a problem if v ends up being in some component C_i later, as we would have a subgraph H_i that is not contained in any region. Thus, we need to somehow preserve edges between $\delta_y(B_i)$ and the residual vertices. We will do this by introducing copies of these edges that we need to preserve, called “zombie edges”. A detailed description of Delete and Separate step as well as the rounding algorithm appears in the next subsection.

3.3 The Delete and Separate Step

We start by describing the Delete and Separate step sketched out in Section 3.2. The main ingredient is the region growing technique.

Region Growing. We begin with some standard definitions. We define the ball $B(s, r)$ centered at vertex s with radius r as $B(s, r) = \{v : d_{sv}^S \leq r\}$. We say that a vertex v is *cut* by $B(s, r)$ if $d_{sv}^S - y_v^S < r < d_{sv}^S$, i.e. v is only “partially inside” the ball, and use $\delta_y(s, r)$ to denote the set of vertices cut by $B(s, r)$; likewise, an edge (u, v) is *cut* by $B(s, r)$ if $d_{su}^S \leq r < d_{su}^S + x_{uv}$ and we use $\delta_x(s, r)$ to denote the set of edges cut by $B(s, r)$. We define two quantities for the x -cost and y -cost. The *cost* LP-cost(s, r) of $B(s, r)$ is defined to be the total cost of the LP solution “contained” in $B(s, r)$:

$$\begin{aligned} \text{LP-cost}(s, r) = & \sum_{u \in B(s, r), v \in B(s, r) \cup \delta_y(s, r)} x_{uv} \\ & + \sum_{(u, v) \in \delta_x(s, r) : u \in B(s, r), v \notin B(s, r)} r - d_{su}^S \end{aligned}$$

Similarly, its *weight* is defined to be the total fractional weight that is “contained” in $B(s, r)$:

$$\text{wt}(s, r) = \sum_{v \in B(s, r)} y_v^S + \sum_{v \in \delta_y(s, r)} r - (d_{sv}^S - y_v^S).$$

See Figure 4 for an illustration. The x -volume and y -volume is then defined to be $\text{vol}_x(s, r) = \frac{\text{LP-cost}(G)}{n^2} + \text{LP-cost}(s, r)$ and $\text{vol}_y(s, r) = \frac{1}{w} + \text{wt}(s, r)$.

These notions also extend to subgraphs. Given a subgraph H' , we define $\text{LP-cost}(H') = \sum_{(u, v) \in H'} x_{uv}$ and $\text{wt}(H') = \sum_{v \in H'} y_v^S$. Similarly, $\text{vol}_x(H') = \frac{\text{LP-cost}(G)}{n^2} + \text{LP-cost}(H')$ and $\text{vol}_y(H') = \frac{1}{w} + \text{wt}(H')$.

The following lemma shows that there exists a *good* radius $r \in [0, 1/12]$ such that the number of edges cut by $B(s, r)$ can be charged to $\text{vol}_x(s, r)$ and simultaneously, the number of vertices cut by $B(s, r)$ can be charged to $\text{vol}_y(s, r)$.




Figure 4: In this figure, the edge lengths represents the x variables, and vertices are represented as circles whose diameters correspond to their y^S weights. The dashed arc is centered at v_1 and has radius 0.8. We have $B(v_1, 0.8) = \{v_1, v_2, v_3\}$. Its boundary edges and vertices are $\delta_x(v_1, 0.8) = \{(v_2, v_4), (v_1, v_4)\}$ and $\delta_y(v_1, 0.8) = \{v_3\}$, respectively. Towards $\text{vol}_x(v_1, 0.8)$, edge (v_1, v_2) contributes 0.1, (v_1, v_3) contributes 0.3, (v_1, v_4) contributes 0.7, and (v_2, v_4) contributes 0.1. Towards $\text{vol}_y(v_1, 0.8)$, v_1 contributes 0.1, v_2 contributes 0.5, and v_3 contributes 0.4.

Lemma 3.3 (Region Growing Lemma). *For every vertex s , we can efficiently find a good radius $r \in [0, 1/12]$ such that*

$$|\delta_x(s, r)| \leq O(\log \log n) \cdot \ln \frac{\text{vol}_x(s, 1/12)}{\text{vol}_x(s, r)} \cdot \text{vol}_x(s, r),$$

$$|\delta_y(s, r)| \leq \kappa \cdot \log w \cdot \text{vol}_y(s, r),$$

for some constant κ .

Proof. Say that r is x -good if it satisfies the first inequality and y -good if it satisfies the second inequality. Set $r(n) := \ln \ln[e(n+1)]$ and define the following sets of radii:

$$A_x = \left\{ r \in [0, 1/12] : \frac{|\delta_x(s, r)|}{\text{vol}_x(s, r)} > 48 \cdot \ln \frac{e \cdot \text{vol}_x(s, 1/12)}{\text{vol}_x(s, r)} \cdot r(n) \right\},$$

$$A_y = \left\{ r \in [0, 1/12] : \frac{|\delta_y(s, r)|}{\text{vol}_y(s, r)} > 48 \cdot \ln(w^2 + 1) \right\}.$$

In other words, A_x and A_y are the sets of radii that are x -bad and y -bad, respectively. We claim that the measure of both A_x and A_y are small: $\mu(A_x), \mu(A_y) \leq 1/48$. The claim then implies that there exists $r \in [0, 1/12]$ that is simultaneously x -good and y -good.

Observe that $\frac{\partial \text{vol}_x(s, r)}{\partial r} = |\delta_x(s, r)|$ and $\frac{\partial \text{vol}_y(s, r)}{\partial r} = |\delta_y(s, r)|$. Suppose, towards a contradiction, that $\mu(A_x) > 1/48$. We have

$$\begin{aligned} & \frac{\partial}{\partial r} \left(-\ln \ln \frac{e \cdot \text{vol}_x(s, 1/12)}{\text{vol}_x(s, r)} \right) \\ &= \frac{|\delta_x(s, r)|}{\text{vol}_x(s, r) \cdot \ln \frac{e \cdot \text{vol}_x(s, 1/12)}{\text{vol}_x(s, r)}} \end{aligned}$$

So

$$\begin{aligned}
& \ln \ln \frac{e \cdot \text{vol}_x(s, 1/12)}{\text{vol}_x(s, 0)} \\
&= - \ln \ln \frac{e \cdot \text{vol}_x(s, 1/12)}{\text{vol}_x(s, r)} \Big|_0^{1/12} \\
&= \int_0^{1/12} \frac{|\delta_x(s, r)|}{\text{vol}_x(s, r) \cdot \ln \frac{e \cdot \text{vol}_x(s, 1/12)}{\text{vol}_x(s, r)}} \partial r \\
&\geq \mu(A_x) \cdot 48 \ln \ln[e(n+1)] > r(n).
\end{aligned}$$

But

$$\ln \ln \frac{e \cdot \text{vol}_x(s, 1/12)}{\text{vol}_x(s, 0)} \leq \ln \ln \frac{e(\text{vol}_x + \text{vol}_x/n)}{\text{vol}_x/n} = r(n).$$

Thus, we have a contradiction and so $\mu(A_x) \leq 1/48$.

We now turn to bounding $\mu(A_y)$. Suppose, towards a contradiction, that $\mu(A_y) > 1/48$.

$$\begin{aligned}
\ln \frac{\text{vol}_y(s, 1/48)}{\text{vol}_y(s, 0)} &= \int_0^{1/12} \frac{1}{\text{vol}_y(s, r)} \frac{\partial \text{vol}_y(s, r)}{\partial r} \partial r \\
&= \int_0^{1/12} \frac{|\delta_y(s, r)|}{\text{vol}_y(s, r)} \partial r \\
&\geq \mu(A_y) \cdot 48 \ln(w^2 + 1) > \ln(w^2 + 1).
\end{aligned}$$

But

$$\ln \frac{\text{vol}_y(s, 1/48)}{\text{vol}_y(s, 0)} \leq \ln \frac{w + 1/w}{1/w} = \ln(w^2 + 1).$$

Therefore, both $\mu(A_x), \mu(A_y) \leq 1/48$ and so there exists $r \in [0, 1/12]$ that is simultaneously x -good and y -good.

To find such an r efficiently, note that as r grows, $\text{vol}_x(s, r)$ and $\text{vol}_y(s, r)$ are non-decreasing. Thus, we only need to check the condition at points r when either $\delta_x(s, r)$ and $\delta_y(s, r)$ changes, which happens at most $2|V'|$ and $2|E'|$ times, respectively. This is because as r grows, once a vertex leaves $\delta_y(s, r)$, it is inside $B(s, r)$ and will never reappear in $\delta_y(s, r)$; and once an edge leaves $\delta_x(s, r)$, both of its endpoints are inside $B(s, r)$ and the edge will never reappear in $\delta_y(s, r)$. \square

The next lemma shows that as we grow a ball $B(s, r)$ around vertex s , as long as $r \leq 1/12$, it cannot contain more than two-thirds of S .

Lemma 3.4. *Let U be a set of vertices. If $\max_{u,v \in U} d_{uv}^S \leq 1/6$, then $|U \cap S| \leq 2|S|/3$.*

Proof. We prove the contrapositive. Suppose, that $|U \cap S| > 2|S|/3$. Let u be a vertex in $U \cap S$. The spreading constraints imply that

$$\sum_{v \in U \cap S} d_{uv}^S \geq |U \cap S| - \frac{|S|}{2} > |U \cap S| - \frac{3|U \cap S|}{4} = \frac{|U \cap S|}{4}.$$

Thus, by averaging, there exists $v \in U \cap S$ such that $d_{uv}^S > 1/4 > 1/6$. \square

We are now ready to describe the PARTITION algorithm, which executes the Delete and Separate component of our algorithm. Recall that in the Delete and Separate step, the input is a subgraph $H = (V', E')$ and a vertex set $S \subseteq V'$, and our goal is find a set D of edges to delete and a vertex set X of size $O(w \log w)$ such that X is a $\frac{2}{3}$ -separator of S in $H - D$.

The Partition algorithm.

1. Initialization. Let $\hat{H} = H$
2. Region growing. While \hat{H} contains more than two-thirds of S ,
 - (a) Find region. Choose an arbitrary vertex $v \in S$ that is contained in \hat{H} and find a good radius r such that $B(v, r)$ satisfies the conditions of Lemma 3.3. Note that distances, balls and boundaries are defined with respect to \hat{H} .
 - (b) Removal. Remove all vertices in $B(v, r)$ and their incident edges from \hat{H} .
 - (c) Add zombies. For each edge (s, u) that was removed, if $s \in \delta_y(v, r)$ and u is still in \hat{H} , add *zombie vertex* $z_u(s)$ with weight $y_{z_u(s)}^S = 0$ to \hat{H} and *zombie edge* $(z_u(s), u)$ with length $x_{z_u(s), u} = x_{su}$.
3. Let $B(v_i, r_i)$ be the ball found in the i -th iteration.
4. Return the vertex set $X = \bigcup_i \delta_y(v_i, r_i)$ and the edge set $D = \bigcup_i \delta_x(v_i, r_i)$, replacing each zombie vertex $z_u(s)$ in X with its original vertex s and each zombie edge $(z_u(s), u)$ in D with its original edge (s, u) .

Lemma 3.5. *Suppose PARTITION took ℓ iterations and let $B(v_1, r_1), \dots, B(v_\ell, r_\ell)$ be the regions it found. Let C_1, \dots, C_k be the components of $H - X - D$. For every $j \in [k]$, define the subgraph $H_j = (V_j, E_j)$ where E_j is the subset of $E' - D$ with at least one endpoint in C_j , and V_j is the set of endpoints of E_j . We have the following:*

1. $|D| \leq O(\log \log n) \cdot \sum_{i=1}^{\ell} \ln \left(\frac{\text{vol}_x(H)}{\text{vol}_x(v_i, r_i)} \right) \text{vol}_x(v_i, r_i)$,
2. $|X| \leq \gamma \cdot (w + |S|/w) \log w$, for some constant γ ,
3. $|C_i \cap S| \leq 2|S|/3$, and
4. *The edge set E_j of each subgraph H_j is either contained in a region $B(v_i, r_i)$ (and so $\text{vol}_x(H_j) \leq \text{vol}_x(v_i, r_i)$) or it is contained in the residual graph \hat{H} at the end of the execution of PARTITION.*

Proof. The first statement follows from the fact that we chose a good radius r_i for each region $B(v_i, r_i)$ and that the x -volume of any region can be at most $\text{vol}_x(H)$. Let us now consider the second statement. The fact that we chose a good radius for each region implies that

$$|X| = \sum_{i=1}^{\ell} |\delta_y(v_i, r_i)| \leq \kappa \cdot \log w \sum_{i=1}^{\ell} \text{vol}_y(v_i, r_i)$$

for some constant κ . Recall that $\text{vol}_y(v_i, r_i) = \text{wt}(v_i, r_i) + 1/w$. Since each vertex can only contribute towards the weight of at most one region and zombie vertices have zero weight, the total weight of any region is at most $\sum_{v \in V'} y_v^S \leq w$. So, $\sum_{i=1}^{\ell} \text{vol}_y(v_i, r_i) \leq (w + \ell/w)$. Moreover, since the center of each region is a vertex v_i of S , we remove at least one vertex of S in each iteration. Thus, the number of iterations ℓ can be at most $|S|$. So, $|X| \leq \kappa \cdot \log w \cdot (w + \ell/w) \leq \kappa \cdot (w + |S|/w) \log w$ and this proves the second statement of the lemma.

Next, we argue that $|C_i \cap S| \leq 2|S|/3$. If C_i was a component that remained at the end of the execution of PARTITION, then by definition, $|C_i \cap S| \leq 2|S|/3$. Suppose C_i was a component that is contained in $B(v_i, r_i)$. By Lemma 3.4, it suffices to check that the distance between any two remaining vertices u and v at the start of iteration i is at least d_{uv}^S . Over the previous iterations,

PARTITION modifies the graph in two ways: by removing edges and vertices, and by introducing zombie edges. Removing edges and vertices clearly cannot decrease the distance between u and v . Zombie edges do not create a shortcut between u and v either. Thus, the distance between u and v in iteration i is at least d_{uv}^S .

Finally, each subgraph H_j is connected (since they are the set of edges with one endpoint in a connected component C_j), thus for each region $B(v_i, r_i)$, either all the vertices of H_j are in $B(v_i, r_i)$ or all of them are not in $B(v_i, r_i)$. \square

3.4 Putting it together

We are now ready to describe the INTERDICT algorithm, which recursively rounds a feasible solution to LP (1). It takes as input a subgraph $H = (V', E')$ and a vertex set S , and deletes a set of edges F such that $\text{tw}(H - F) \leq O(w \log w)$. Let γ be the constant in Lemma 3.5.

The Interdict Algorithm.

1. Delete and Separate. Use algorithm PARTITION to find a set D of edges to delete and a $\frac{2}{3}$ -separator X of S in $H - D$. Let C_1, \dots, C_k be the components of $H - D - X$. Delete D .
2. Define subproblems. For $i \in [k]$, define the subgraph $H_i = (V_i, E_i)$ where E_i is the subset of $E' - D$ with an endpoint in C_i , and V_i is the set of endpoints of E_i .
3. Recurse. For $i \in [k]$, call INTERDICT($H_i, V_i \cap (X \cup S)$).

To round a feasible solution of LP (1), we call INTERDICT(G, S_0) where S_0 is an arbitrary set of at most $O(w \log w)$ vertices. Let F be the set of edges deleted by INTERDICT(G, S_0).

Lemma 3.6. $\text{tw}(G - F) \leq O(w \log w)$.

Proof. We can apply exactly the same analysis of the classic tree decomposition algorithm [37] to argue that the $O(w \log w)$ -size separators found in the recursion can be used to construct a tree decomposition T of $G - F$ such that the width of T is at most $O(w \log w)$. \square

Recall that $\text{LP-cost}(G) = \sum_{(u,v) \in E} x_{uv}$, the cost of the LP solution $(\mathbf{x}, \mathbf{y}, \mathbf{d})$.

Lemma 3.7. $|F| \leq O(\log n \log \log n) \text{LP-cost}(G)$.

Proof. Let k be the recursion depth of INTERDICT(G, S_0). For each depth j , let \mathcal{H}_j be the collection of subgraphs that INTERDICT recursed on at depth j , and for each $H \in \mathcal{H}_j$, let $R(H)$ be the set of regions found by PARTITION when INTERDICT recursed on H . (Note that at depth 0, we have $\mathcal{H}_0 = \{G\}$.) By Lemma 3.5, we have the following bound on $|F|$:

$$|F| \leq O(\log \log n) \cdot \sum_j \sum_{H \in \mathcal{H}_j} \sum_{B(s,r) \in R(H)} \ln \left(\frac{\text{vol}_x(H)}{\text{vol}_x(s,r)} \right) \text{vol}_x(s,r). \quad (2)$$

For each edge $(u, v) \in E$, define $g(u, v) = x_{uv} + \frac{\text{LP-cost}(G)}{n^2}$. We have that

$$\sum_{(u,v) \in E} g(u, v) \leq 2 \text{LP-cost}(G). \quad (3)$$

We now show that $\sum_j \sum_{H \in \mathcal{H}_j} \sum_{B(s,r) \in R(H)} \ln \left(\frac{\text{vol}_x(H)}{\text{vol}_x(s,r)} \right) \text{vol}_x(s,r) \leq O(\log n) \sum_{(u,v) \in E} g(u, v)$. We will do this by charging $\sum_{H \in \mathcal{H}_j} \sum_{B(s,r) \in R(H)} \ln \left(\frac{\text{vol}_x(H)}{\text{vol}_x(s,r)} \right) \text{vol}_x(s,r)$ to edges of the subgraphs

in \mathcal{H}_j and proving that every edge (u, v) receives a total charge, across all depths, of at most $O(\log n)g(u, v)$. Our charging scheme is as follows: for each subgraph $H \in \mathcal{H}_j$ and each region $B(s, r) \in R(H)$, charge $\ln \left(\frac{\text{vol}_x(H)}{\text{vol}_x(s, r)} \right) g(u, v)$ to each edge (u, v) with at least one endpoint in $B(s, r)$.

Let us first see why the total amount charged per region $B(s, r) \in R(H)$ is at least $\ln \left(\frac{\text{vol}_x(H)}{\text{vol}_x(s, r)} \right) \text{vol}_x(s, r)$. Let $E_H(s, r)$ be the edges with at least one endpoint in $B(s, r)$. This is exactly the set of edges that contribute to the region's x -volume $\text{vol}_x(s, r)$. Therefore,

$$\begin{aligned} \sum_{(u,v) \in E_H(s,r)} g(u, v) &\geq \sum_{(u,v) \in E_H(s,r)} x_{uv} + \frac{\text{LP-cost}(G)}{n^2} \\ &\geq \text{vol}_x(s, r) \end{aligned}$$

and so the total amount charged is at least $\ln \left(\frac{\text{vol}_x(H)}{\text{vol}_x(s, r)} \right) \text{vol}_x(s, r)$. Overall, we have

$$\sum_j \sum_{H \in \mathcal{H}_j} \sum_{B(s,r) \in R(H)} \ln \left(\frac{\text{vol}_x(H)}{\text{vol}_x(s, r)} \right) \text{vol}_x(s, r) \leq \sum_{(u,v) \in E} \phi(u, v), \quad (4)$$

where $\phi(u, v)$ is the total charge received by (u, v) .

Fix an edge $(u, v) \in E$. Let us now upper bound $\phi(u, v)$. Since INTERDICT recurses on edge-disjoint subgraphs, (u, v) is contained in at most one subgraph of \mathcal{H}_j for each depth j , and if it does not belong to a subgraph of \mathcal{H}_j , then it does not belong to any subgraph of $\mathcal{H}_{j'}$ at lower depths $j' > j$. In the worst case, (u, v) is contained in some subgraph $H_j \in \mathcal{H}_j$ for every depth j . Let us assume this is so and define $B(s_j, r_j) \in R(H_j)$ to be the region containing (u, v) . Lemma 3.5 tells us that all of H_j is contained in the region $B(s_{j-1}, r_{j-1})$ for each j and so $\text{vol}_x(s_j, r_j) \geq \text{vol}_x(H_{j+1})$. Thus,

$$\begin{aligned} \phi(u, v) &= \sum_{j=0}^k \ln \left(\frac{\text{vol}_x(H_j)}{\text{vol}_x(s_j, r_j)} \right) g(u, v) \\ &\leq \sum_{j=0}^k \ln \left(\frac{\text{vol}_x(H_j)}{\text{vol}_x(H_{j+1})} \right) g(u, v) \\ &\leq \ln \left(\frac{\text{vol}_x(G)}{\text{vol}_x(H_{k+1})} \right) g(u, v) \\ &\leq \ln(n^2 + 1)g(u, v), \end{aligned}$$

where the last inequality follows from the fact that x -volume is always at least $\frac{\text{LP-cost}(G)}{n^2}$.

Combining this with Inequality (4), we get

$$\begin{aligned} &\sum_j \sum_{H \in \mathcal{H}_j} \sum_{B(s,r) \in R(H)} \ln \left(\frac{\text{vol}_x(H)}{\text{vol}_x(s, r)} \right) \text{vol}_x(s, r) \\ &\leq \sum_{(u,v) \in E} \ln(n^2 + 1)g(u, v) \\ &\leq 2 \ln(n^2 + 1) \text{LP-cost}(G), \end{aligned}$$

where the last inequality follows from Inequality (3). Finally, plugging this into the right hand side of (2) gives us $|F| \leq O(\log n \log \log n) \text{LP-cost}(G)$, as desired. \square

Lemmas 3.7 and 3.6 imply Lemma 3.2.

4 Using the LP

We now come to the problem of how to handle the LP (1) in $n^{O(1)}$ time. As discussed in the Introduction, we bypass the need to completely solve this LP using the Round-or-Separate framework. A crucial ingredient here is that if the rounding step gets stuck (is unable to find a small separator), we need to find a violated inequality for the x -variables, and not just for the LP (1). In Section 4.1 we show how to reformulate LP (1) to only have the x variables. In this reformulation, the coefficients of the inequalities come from feasible points in a polytope with exponentially many variables, so Section 4.2 deals with how to fix such a violated inequality.

4.1 Reformulating the LP

We first reformulate LP (1)—the original LP with variables $(\mathbf{x}, \mathbf{y}, \mathbf{d})$ —to obtain another LP with only x -variables.

Call a vector \mathbf{x} *feasible* if there exist vectors \mathbf{y} and \mathbf{d} such that $(\mathbf{x}, \mathbf{y}, \mathbf{d})$ is a feasible solution to LP (1). Define \mathcal{F} to be the set of all feasible vectors \mathbf{x} , and observe that \mathcal{F} is simply the feasible region of LP (1) with the y and d variables projected out.

Next, we show how to describe \mathcal{F} using linear inequalities. For every vertex set S , we define an LP parameterized by \mathbf{x} (see Figure 5). We call this sep-LP(\mathbf{x}, S). We emphasize that in this LP,

$$\begin{array}{ll}
 \min & \sum_v y_v \\
 \text{s.t.} & d_{uv} \leq \sum_{e \in E(P)} x_e + \sum_{t \in V(P)} y_t \quad \forall u, v \in V, P \in \mathcal{P}(u, v) \\
 & \sum_{u \in T} d_{uv} \geq |T| - |S|/2 \quad \forall T \subseteq S \subseteq V, v \in T
 \end{array} \tag{5}$$

Figure 5: sep-LP(\mathbf{x}, S)

only \mathbf{y} and \mathbf{d} are variables. Recall that this LP is related to the problem of finding a small balanced vertex separator of S in G , provided the edges are removed fractionally to extent x_e .

Definition 4.1. Let $\lambda(\mathbf{x}, S)$ be the value of the optimum solution to the sep-LP(\mathbf{x}, S). We say that \mathbf{x} is S -feasible if $\lambda(\mathbf{x}, S) \leq w$, and denote by $\mathcal{F}(S)$ the set of S -feasible vectors.

Lemma 4.2. $\mathcal{F} = \bigcap_{S \subseteq V} \mathcal{F}(S)$.

Proof. Suppose $\mathbf{x} \in \mathcal{F}$. Then, for every vertex set S , the vectors \mathbf{y}^S and \mathbf{d}^S are a feasible solution to sep-LP(\mathbf{x}, S) and $\sum_v y^S \leq w$, so $\lambda(\mathbf{x}, S) \leq w$.

On the other hand, suppose $\lambda(\mathbf{x}, S) \leq w$ for all vertex sets S . Define \mathbf{y} and \mathbf{d} such that for each $S \in \mathcal{S}$, $(\mathbf{y}^S, \mathbf{d}^S)$ is the optimal solution to sep-LP(\mathbf{x}, S). As $\lambda(\mathbf{x}, S) \leq w$ for all $S \in \mathcal{S}$, we have that $(\mathbf{x}, \mathbf{y}, \mathbf{d})$ is a feasible solution to LP (1). \square

Thus, to describe \mathcal{F} using linear inequalities, it suffices to describe $\mathcal{F}(S)$ using linear inequalities. By linear programming duality, we have that $\lambda(\mathbf{x}, S) \leq w$ if and only if for every feasible solution to the dual of sep-LP(\mathbf{x}, S) has objective value at most w .

So we consider the dual LP (6) (Figure 6), and denote it by flow-LP(\mathbf{x}, S). (We call this flow-LP as it is dual to a “cut-type” LP.)

$$\begin{array}{ll}
\max & \sum_{T \subseteq S, v \in T} g_{T,v} (|T| - |S|/2) - \sum_{u,v \in V; P \in \mathcal{P}(u,v)} f_P^{uv} \left[\sum_{e \in P} x_e \right] \\
\text{s.t.} & \sum_{T \subseteq V: u,v \in T} g_{T,u} + g_{T,v} \leq \sum_{P \in \mathcal{P}(u,v)} f_P^{uv} \quad \forall u, v \in V, \\
& \sum_{u,v \in V} \sum_{P \in \mathcal{P}(u,v): t \in V(P)} f_P^{uv} \leq 1 \quad \forall t \in V
\end{array} \tag{6}$$

Figure 6: flow-LP(\mathbf{x}, S).

Lemma 4.3. *The dual to sep-LP(\mathbf{x}, S) is given by LP (6) below.*

Proof. Let us introduce the variables f_P^{uv} for the first set of constraints in LP (5), and the variables $g_{T,v}$ for the second set of constraints. Again, we emphasize that only f and g are variables here (and that x is not a variable).

Let us check that the dual is exactly LP (6). Rewrite the first set of primal constraints as

$$\sum_{t \in V(P)} y_t - d_{uv} \geq - \sum_{e \in E(P)} x_e \quad \forall u, v \in V, P \in \mathcal{P}(u, v)$$

In the objective, the coefficient of f_P^{uv} is $-\sum_{e \in E(P)} x_e$ and the coefficient of $g_{T,v}$ is $|T| - |S|/2$. The dual constraints correspond to primal variables d_{uv} and y_t .

In the primal, the variable d_{uv} has a coefficient of -1 for the constraints corresponding to the dual variables f_P^{uv} for $P \in \mathcal{P}(u, v)$, and a coefficient of 1 in the constraints corresponding to the dual variables $g_{T,t}$ for $t \in \{u, v\}$ and $T \ni t$. Similarly, the variable y_t has a coefficient of 1 for the constraints corresponding to the dual variables f_P^{uv} for $u, v \in V$ and $P \in \mathcal{P}(u, v)$ such that $t \in V(P)$. So, LP (6) is dual to LP (5). \square

Separation oracle for $\mathcal{F}(S)$. Given a solution (\mathbf{f}, \mathbf{g}) to flow-LP(\mathbf{x}, S), we denote its objective value as flow-LP($\mathbf{x}, S, \mathbf{f}, \mathbf{g}$). LP duality implies the following lemma.

Lemma 4.4. *$\mathbf{x} \in \mathcal{F}(S)$ if and only if flow-LP($\mathbf{x}, S, \mathbf{f}, \mathbf{g}$) $\leq w$ for all feasible dual solutions (\mathbf{f}, \mathbf{g}) .*

Let us see what we have achieved so far. The expression flow-LP($\mathbf{x}, S, \mathbf{f}, \mathbf{g}$) $\leq w$ is a linear constraint on x , whose coefficients (\mathbf{f}, \mathbf{g}) are feasible points in the polytope given by (6). Now we crucially note that the coefficients of the constraints on (\mathbf{f}, \mathbf{g}) depend only on the topology of the graph G and not on \mathbf{x} .

Definition 4.5. *We say that (\mathbf{f}, \mathbf{g}) is S -valid if it satisfies the constraints of flow-LP(\mathbf{x}, S).*

Thus, Lemma 4.4 gives a description of $\mathcal{F}(S)$ in terms of linear inequalities. Together with Lemma 4.2, we get a description of \mathcal{F} in terms of linear inequalities and so we can reformulate LP (1) as follows.

$$\begin{array}{ll}
\min & \sum_{(u,v) \in E} x_{uv} \\
\text{s.t.} & \text{flow-LP}(\mathbf{x}, S, \mathbf{f}, \mathbf{g}) \leq w \quad \forall S \subseteq V, (\mathbf{f}, \mathbf{g}) \text{ } S\text{-valid}
\end{array} \tag{7}$$

4.2 Round or Separate

We will work with the reformulated LP (7). One immediate obstacle is that it is unclear how to get an efficient separation oracle for \mathcal{F} . In fact, even for a fixed vertex set S , it is not immediately clear how to find a violated inequality for $\mathcal{F}(S)$. The problem is that for any constraint, the coefficients of the x -variables are based on \mathbf{f}, \mathbf{g} , which need to satisfy LP (6), and it is not clear how to generate them; e.g. there are exponentially many \mathbf{f}, \mathbf{g} variables in flow-LP(\mathbf{x}, S). To get over these problems, we do the following.

We apply the Round-or-Separate framework with the INTERDICT algorithm in Section 3.1. Roughly speaking, we start with some candidate solution \mathbf{x} (possibly infeasible), and try to construct a tree decomposition of width $O(w \log w)$ using INTERDICT. If we succeed, we are done. Otherwise, INTERDICT could not find a small balanced separator for some vertex set S , and this can only happen if $\mathbf{x} \notin \mathcal{F}(S)$. Later, in Lemma 4.6, we give an efficient procedure that given \mathbf{x} and S , determines whether $\mathbf{x} \in \mathcal{F}(S)$ and if so, outputs a solution $(\mathbf{y}^S, \mathbf{d}^S)$ that is feasible to sep-LP(\mathbf{x}, S) and satisfies $\sum_v y_v^S \leq w$; otherwise, it outputs a separating hyperplane, i.e. an S -valid (\mathbf{f}, \mathbf{g}) such that flow-LP($\mathbf{x}, S, \mathbf{f}, \mathbf{g}$) $> w$. In the first case, INTERDICT can use $(\mathbf{y}^S, \mathbf{d}^S)$ to find a $O(w \log w)$ -size separator of S and make progress. In the second case, we can add the separating hyperplane to find a new candidate x and repeat the whole tree decomposition procedure. By the standard Ellipsoid method-based, separation-versus-optimization framework, the number of such iterations is polynomially bounded.

It remains to show how to efficiently separate $\mathcal{F}(S)$ for any given subset S . We will do this indirectly by using the Ellipsoid Method and applying duality to sep-LP(\mathbf{x}, S).

Lemma 4.6. *There exists a polynomial-time algorithm that given \mathbf{x} and S , determines whether $\mathbf{x} \in \mathcal{F}(S)$, and if so, outputs a solution $(\mathbf{y}^*, \mathbf{d}^*)$ that is feasible to sep-LP(\mathbf{x}, S) and satisfies $\sum_v y_v^* \leq w$; otherwise, it finds a violated inequality for LP (7).*

Proof. We apply the Ellipsoid method to find an optimal solution $(\mathbf{y}^*, \mathbf{d}^*)$ to sep-LP(\mathbf{x}, S). We can do this efficiently by using a separation oracle that uses Dijkstra's algorithm to determine the distances \mathbf{d} given \mathbf{x} and \mathbf{y} . If $\sum_v y_v^* \leq w$, then $\mathbf{x} \in \mathcal{F}(S)$ by definition.

On the other hand, if $\sum_v y_v^* > w + \epsilon$ (where ϵ can be made exponentially small), then we can find a violated inequality for (7) as follows. As the Ellipsoid method added only polynomially many constraints we can re-solve this LP only on these added constraints and assume that \mathbf{y}^* is a solution to this smaller LP on only polynomially many constraints. By complementary slackness, there exists an optimal dual solution $(\mathbf{f}^*, \mathbf{g}^*)$ where the only non-zero dual variables are those that correspond to these polynomially many primal constraints. Thus, it suffices to solve flow-LP(\mathbf{x}, S) restricted to these dual variables which makes flow-LP(\mathbf{x}, S) polynomial in size. The objective function in flow-LP then gives the violated inequality for LP (7). \square

5 Bounded Size Interdiction and MIS in Noisy Planar Graphs

We consider the Bounded Size Interdiction problem and show how it implies Theorem 1.2.

Consider the noisy graph G obtained by adding δn edges to some planar G_0 . Let us view G as obtained by superimposing the noisy edges on the recursive decomposition of G_0 given by Lemma 2.1. This directly implies the following (noisy) decomposition for G .

Lemma 5.1 (Noisy Decomposition). *Given a δ -noisy planar graph G , for any $\alpha > 0$, there exists a partition X, C_1, \dots, C_k of V with (i) $|X| \leq c\alpha n$ for some universal constant c , (ii) $|C_i| \leq 1/\alpha^2$ for all $i \in [k]$, and (iii) at most δn edges whose endpoints lie in distinct (two different) C_i s.*

Of course as we do not know G_0 , it is not clear how to find such a decomposition. Theorem 5.2 shows that this can be done approximately.

Theorem 5.2 (Noisy Bounded Size Interdiction). *Let $G = (V, E)$ be any graph that has a vertex partition X, C_1, \dots, C_k with $|C_i| \leq s$ for each $i \in [k]$, and let b be the total number of edges whose endpoints lie in distinct C_i s.*

Then for every $\beta \leq 1$, we can find in time $n^{O(s)}$ a vertex partition $X', C'_1, \dots, C'_{k'}$ such that (i) $|X'| = O(|X|/\beta)$, (ii) $|C'_i| \leq s$ for each $i \in [k']$, and (iii) at most $O(b + \beta|E|)$ edges whose endpoints lie in distinct C'_i s.

This implies the following proper decomposition (where the C'_i are components of $G[V - X']$).

Corollary 5.3. *There is a subset $X' \subset V$ of size $O(|X|/\beta + b + \beta|E|)$ such that the components in $G[V - X']$ have size at most s .*

Proof. For each edge (u, v) with $u \in C'_i$ and $v \in C'_j$ for $i \neq j \in [k']$, put an endpoint (say u) in X' and remove u from C'_i . As a result, X' has size $O(|X|/\beta + b + \beta|E|)$ and there are no edges between different C'_i and C'_j . \square

Before we prove Theorem 5.2, let us first see how it implies Theorem 1.2.

Proof of Theorem 1.2. Given an $\epsilon > 0$, set $\gamma = O(\epsilon^2)$. Applying Lemma 5.1 with $\alpha = \gamma$, G has a noisy decomposition with $|X| = O(\gamma n)$ and pieces C_i of size $s = 1/\gamma^2$ and at most $b = \delta n$ edges between these C_i 's. Moreover G has at most $(3 + \delta)n \leq 4n$ edges. Applying Corollary 5.3 with $\beta = \gamma^{1/2}$, gives an X' of size $O(|X|/\gamma^{1/2} + \delta n + \gamma^{1/2}n) = O((\epsilon + \delta)n)$, such that the components of $G[V - X']$ have at most $1/\gamma^2 = 1/\epsilon^4$ vertices.

By exhaustive search, find an MIS in each C_i separately and return their union I . Clearly, I is a valid independent set. By Theorem 5.2, the algorithm has overall running time $n^{O(1/\gamma^2)} = n^{O(1/\epsilon^4)}$. As G_0 is planar and hence 4-colorable, and $G \setminus G_0$ has at most δn noisy edges, $\alpha(G) \geq \alpha(G_0) - \delta n \geq n/4 - \delta n$. Moreover, as $|I| \geq \alpha(G) - |X'|$, this gives $|I| \geq (1 - O(\epsilon + \delta))\alpha(G)$. \square

Remark. Theorem 1.2 extends directly to the minor-free case using the separator theorem for minor-free graphs [1], and the fact that these graphs have bounded average degree and thus $\alpha(G) = \Omega(n)$.

5.1 Proof of Theorem 5.2

LP formulation. Given $0 < \beta \leq 1$ and G as input, we first write an integer program to find X and the C_i s. Let $a = |X|$ (we can assume that a is known as the algorithm can try every value). For each vertex v , the variable y_v indicates if $v \in X$. For each subset $S \subset V$ with $|S| \leq s$, the variable z_S indicates if S is one of the pieces C_i . Let \mathcal{S} be the collection of all such subsets S . For each $(u, v) \in E$, the variable x_{uv} indicates whether the edge (u, v) is such that $u \in C_i$ and $v \in C_j$ for some $i \neq j$.

Consider the integer program IP (8) in Figure 7. This is easily seen to be a valid formulation for the problem. The first set of constraints ensure that X contains at most a vertices. The second set of constraints ensure that each vertex lies in either X or some C_i . The third set of constraints are more involved and force x_{uv} to be 1 if some edge has endpoints in distinct C_i and C_j . In particular, it says that if u lies in some C_i and v does not lie in that C_i , then either v lies in X (i.e. $y_v = 1$) or $x_{uv} = 1$. Note that the third set of constraints are asymmetric in u and v , and we will put two such constraints (with u and v swapped) for each edge (u, v) .

As the objective function exactly measures the number of edges with endpoints in two different C_i 's, it follows that the IP above has a feasible solution with value at most δn .

$$\begin{array}{ll}
\min & \sum_{(u,v) \in E} x_{uv} \\
\text{s.t.} & \sum_v y_v \leq a \\
& \sum_{S: v \in S} z_S = 1 - y_v \quad \forall v \in V \\
& \sum_{S: u \in S \wedge v \notin S} z_S \leq x_{uv} + y_v \quad \forall (u,v) \in E \\
& x_{uv}, y_v, z_S \in \{0, 1\} \quad \forall (u,v) \in E, v \in V, S \in \mathcal{S}
\end{array} \tag{8}$$

Figure 7: IP for the Bounded Size Interdiction Problem.

Consider the LP relaxation of this program. It has $O(n^s)$ variables, and $O(n^2)$ non-trivial constraints. So in time $n^{O(s)}$, we can find some basic feasible solution with support size at most $O(n^2)$. We reuse x_{uv} , y_v and z_S to denote some fixed optimum solution to the LP.

The Rounding Algorithm. The algorithm will construct the required X' and the collection \mathcal{C} of sets C'_i from the LP solution by the following preprocessing and sampling procedure.

1. Initialization. We set $X', \mathcal{C} = \emptyset$ and $U = V$, where U denotes the set of vertices not covered by \mathcal{C} .
2. Preprocessing. Add every vertex v with $y_v \geq \beta$ to X' . Set $U = U \setminus X'$
3. Sampling to create \mathcal{C} . Arbitrarily order the sets S_1, \dots, S_k in the support of the LP solution. Repeat the following (phase) until U is empty:
Phase. For $i = 1, \dots, k$, sample the set S_i randomly with probability z_{S_i} . If S_i is picked, add $C' = S_i \cap U$ to the collection \mathcal{C} , and update $U = U \setminus S_i$.

Analysis. Clearly the sets C' produced by the algorithm have size at most s , and they are disjoint.

Lemma 5.4. $|X'| \leq a/\beta$.

Proof. As X' is the set of vertices v with $y_v \geq \beta$, there can be at most a/β such vertices by the LP constraint $\sum_v y_v \leq a$. \square

Henceforth, we also assume that $\beta \leq 1/2$, otherwise choosing $X' = \emptyset$ and partitioning V arbitrarily into sets C_1, \dots, C_k of size at most s trivially suffices for Theorem 5.2.

We now show that the algorithm runs in expected polynomial time and does not generate a vertex partition with too many edges between distinct C'_i .

Lemma 5.5. *Let F be the set of edges with endpoints in two distinct C'_i s. Then $\mathbb{E}[|F|] \leq O(b + \beta|E|)$. Moreover, the algorithm terminates after at most $O(n^2 \log n)$ sampling steps with high probability.*

Proof. We claim that U is empty after $O(\log n)$ phases, with high probability. After the preprocessing step, each uncovered vertex in U has $y_v < \beta \leq 1/2$. Thus, by the second LP constraint, $p_v := \sum_{S \ni v} z_S = 1 - y_v \geq 1/2$. So the probability that a vertex v is not covered after j phases is

$$\left(\prod_{S \ni v} (1 - z_S) \right)^j \leq \exp(-jp_v) \leq \exp(-j/2) \leq (2/3)^j$$

The claim now follows from a union bound over the n vertices.

We now bound the size of F . Let us focus on an edge $e = (u, v)$ and bound the probability that it is *cut*, that is, added to F during the Sampling step. Let U_j denote the vertices in U at the end of phase j . The edge is cut in phase j if and only if both u and v remain in U at the end of phase $j - 1$ (i.e. $u, v \in U_{j-1}$) and a set S with $|S \cap \{u, v\}| = 1$ is chosen in phase j . As $\Pr[u, v \in U_{j-1}] \leq \Pr[v \in U_{j-1}] \leq (2/3)^{j-1}$, this implies that

$$\Pr[(u, v) \text{ cut in phase } j] \leq (2/3)^{j-1} \cdot \sum_{S: |S \cap \{u, v\}|=1} z_S. \quad (9)$$

Moreover, by the third set of constraints in LP (8)

$$\sum_{S: |S \cap \{u, v\}|=1} z_S = \sum_{S: u \in S \wedge v \notin S} z_S + \sum_{S: v \in S \wedge u \notin S} z_S \leq 2x_{uv} + y_u + y_v. \quad (10)$$

Summing (9) over all the phases and using (10), we get

$$\Pr[(u, v) \text{ cut}] \leq 3(2x_{uv} + y_u + y_v) \leq 6x_{uv} + 6\beta,$$

where the second inequality follows as both u and v were not chosen in X during the preprocessing step and hence $y_u, y_v \leq \beta$. By linearity of expectation, this implies that

$$\mathbb{E}[|F|] = \sum_{(u, v) \in E} (6x_{uv} + 6\beta) = O(b + \beta|E|).$$

□

Acknowledgements

We would like to thank Marcin Pilipczuk and Fedor Fomin for many useful discussions and pointers, and anonymous reviewers for bringing recent local search-based techniques to our attention. Part of this work was done when all the authors were visiting the Simons Institute for the Theory of Computing.

References

- [1] Noga Alon, Paul Seymour, and Robin Thomas. A separator theorem for nonplanar graphs. *Journal of the American Mathematical Society*, 3(4):801–808, 1990.
- [2] Eyal Amir. Approximation algorithms for treewidth. *Algorithmica*, 56(4):448–479, 2010.
- [3] Hyung-Chan An, Mohit Singh, and Ola Svensson. Lp-based algorithms for capacitated facility location. In *Foundations of Computer Science, FOCS*, pages 256–265, 2014.

- [4] Sanjeev Arora, Satish Rao, and Umesh V. Vazirani. Expander flows, geometric embeddings and graph partitioning. *J. ACM*, 56(2), 2009.
- [5] Brenda S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *Journal of the ACM (JACM)*, 41(1):153–180, 1994.
- [6] Hans L. Bodlaender. Dynamic programming on graphs with bounded treewidth. In *ICALP*, pages 105–118, 1988.
- [7] Hans L. Bodlaender, John R. Gilbert, Hjálmtýr Hafsteinsson, and Ton Kloks. Approximating treewidth, pathwidth, frontsize, and shortest elimination tree. *J. Algorithms*, 18(2):238–255, 1995.
- [8] Timothy M. Chan and Sarel Har-Peled. Approximation algorithms for maximum independent set of pseudo-disks. *Discrete & Computational Geometry*, 48(2):373–392, 2012.
- [9] Chandra Chekuri and Anastasios Sidiropoulos. Approximation algorithms for euler genus and related problems. In *Foundations of Computer Science, FOCS*, pages 167–176, 2013.
- [10] Julia Chuzhoy, Yury Makarychev, and Anastasios Sidiropoulos. On graph crossing number and edge planarization. In *SODA 2011*, pages 1050–1069, 2011.
- [11] Bruno Courcelle and Joost Engelfriet. *Graph Structure and Monadic Second-Order Logic - A Language-Theoretic Approach*, volume 138 of *Encyclopedia of mathematics and its applications*. Cambridge University Press, 2012.
- [12] Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015.
- [13] D. Erik Demaine and MohammadTaghi Hajiaghayi. *Encyclopedia of Algorithms*, chapter Approximation Schemes for Planar Graph Problems, pages 59–62. Springer US, 2008.
- [14] Erik D. Demaine and MohammadTaghi Hajiaghayi. The bidimensionality theory and its algorithmic applications. *Comput. J.*, 51(3):292–302, 2008.
- [15] Erik D. Demaine, MohammadTaghi Hajiaghayi, and Ken-ichi Kawarabayashi. Algorithmic graph minor theory: Decomposition, approximation, and coloring. In *FOCS 2005*, pages 637–646, 2005.
- [16] Matt DeVos, Guoli Ding, Bogdan Oporowski, Daniel P. Sanders, Bruce Reed, Paul Seymour, and Dirk Vertigan. Excluding any graph as a minor allows a low tree-width 2-coloring. *Journal of Combinatorial Theory, Series B*, 91(1):25–41, 2004.
- [17] David Eppstein. Diameter and treewidth in minor-closed graph families. *Algorithmica*, 27(3-4):275–291, 2000.
- [18] David Eppstein. Graph families. *Algorithmica*, 27(3):275–291, 2000.
- [19] Guy Even, Joseph (Seffi) Naor, Satish Rao, and Baruch Schieber. Divide-and-conquer approximation algorithms via spreading metrics. *Journal of the ACM (JACM)*, 47(4):585–616, 2000.
- [20] Uriel Feige, MohammadTaghi Hajiaghayi, and James R. Lee. Improved approximation algorithms for minimum weight vertex separators. *SIAM J. Comput.*, 38(2):629–657, 2008.

- [21] Uriel Feige and Joe Kilian. Heuristics for finding large independent sets, with applications to coloring semi-random graphs. In *Foundations of Computer Science, FOCS*, pages 674–683, 1998.
- [22] Fedor V. Fomin, Daniel Lokshtanov, Neeldhara Misra, and Saket Saurabh. Planar f -deletion: Approximation, kernelization and optimal FPT algorithms. In *Foundations of Computer Science, FOCS*, pages 470–479, 2012.
- [23] Fedor V Fomin, Daniel Lokshtanov, Venkatesh Raman, and Saket Saurabh. Bidimensionality and eptas. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms*, pages 748–759. SIAM, 2011.
- [24] Greg N. Frederickson. Fast algorithms for shortest paths in planar graphs, with applications. *SIAM J. Comput.*, 16(6):1004–1022, 1987.
- [25] Eugene C Freuder. Complexity of k -tree structured constraint satisfaction problems. In *Proceedings of the eighth National conference on Artificial intelligence-Volume 1*, pages 4–9. AAAI Press, 1990.
- [26] Avinatan Hassidim, Jonathan A. Kelner, Huy N. Nguyen, and Krzysztof Onak. Local graph partitions for approximation and testing. In *Foundations of Computer Science, FOCS*, pages 22–31, 2009.
- [27] Johan Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. In *Foundations of Computer Science, FOCS*, pages 627–636, 1996.
- [28] Johan Håstad. Some optimal inapproximability results. *Journal of the ACM (JACM)*, 48(4):798–859, 2001.
- [29] Sanjeev Khanna and Rajeev Motwani. Towards a syntactic characterization of ptas. In *Symposium on Theory of Computing, STOC '96*, pages 329–337, 1996.
- [30] Euiwoong Lee. Partitioning a graph into small pieces with applications to path transversal. In *SODA 2017*, pages 1546–1558, 2017.
- [31] Retsef Levi, Andrea Lodi, and Maxim Sviridenko. Approximation algorithms for the capacitated multi-item lot-sizing problem via flow-cover inequalities. *Math. Oper. Res.*, 33(2):461–474, 2008.
- [32] Richard J. Lipton and Robert E. Tarjan. A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics*, 36(2):177–189, 1979.
- [33] Avner Magen and Mohammad Moharrami. Robust algorithms for max independent set on minor-free graphs based on the sherali-adams hierarchy. In *APPROX*, pages 258–271, 2009.
- [34] Konstantin Makarychev, Yury Makarychev, and Aravindan Vijayaraghavan. Approximation algorithms for semi-random partitioning problems. In *Symposium on Theory of Computing, STOC*, pages 367–384, 2012.
- [35] Dániel Marx. Can you beat treewidth? In *Foundations of Computer Science, 2007. FOCS'07. 48th Annual IEEE Symposium on*, pages 169–179. IEEE, 2007.

- [36] Bruce A Reed. Tree width and tangles: A new connectivity measure and some applications. In R. A. Bailey, editor, *Surveys in Combinatorics, 1997*, pages 87–162. Cambridge University Press, 1997. Cambridge Books Online.
- [37] Neil Robertson and Paul D. Seymour. Graph minors .xiii. the disjoint paths problem. *J. Comb. Theory, Ser. B*, 63(1):65–110, 1995.
- [38] Paul D. Seymour. Packing directed circuits fractionally. *Combinatorica*, 15(2):281–288, 1995.
- [39] Daniel A. Spielman and Shang-Hua Teng. Smoothed analysis of algorithms: Why the simplex algorithm usually takes polynomial time. *J. ACM*, 51(3):385–463, 2004.
- [40] Yu Wu, Per Austrin, Toniann Pitassi, and David Liu. Inapproximability of treewidth and related problems. *J. Artif. Intell. Res. (JAIR)*, 49:569–600, 2014.